

WHAT IS CLAIMED IS:

1. A method for conducting physical address discovery, facilitating point-to-point communications between hosts of a cluster operating in a cluster mode wherein acceptable messages are addressed to a shared cluster address, the method comprising the
5 steps of:

receiving an address discovery request initiated by a source host seeking a physical address of a target host within the cluster; and

generating an address discovery response acceptable by the source host operating in the cluster mode and including:

10 a response source physical address field specifying a non-cluster mode physical address of the target host.

2. The method of claim 1 wherein the address discovery request is an ARP request.

15 3. The method of claim 1 wherein the non-cluster mode physical address of the target host is a dedicated address of the target host.

4. The method of claim 3 wherein the dedicated address is derived from an
20 IP address assigned to the target host.

5. The method of claim 1 wherein the non-cluster mode physical address of the target host is a shared address assigned to multiple hosts within the cluster.

25 6. The method of claim 1 wherein the generating step is executed in accordance with a further step of:

determining whether the address discovery request was issued by a source host operating in the cluster mode.

7. The method of claim 6 wherein the receiving, determining and generating steps are performed on a responding host that is distinct from an initiating host from which the address discovery request originates.

5 8. The method of claim 7 wherein the determining step comprises detecting that the address discovery request includes:

a request source physical address field specifying the shared cluster address assigned to the cluster; and
a request source network communication protocol-specific address field identifying a host within the cluster of hosts.

10 9. The method of claim 8 wherein the shared cluster address is a MAC address.

15 10. The method of claim 9 wherein the source network communication protocol-specific address field contains an IP address.

11. The method of claim 8 further comprising the step of:
maintaining, by the responding host, a list of network communication protocol-specific addresses corresponding to hosts within the cluster.

20 12. The method of claim 8 further comprising the steps of:
modifying the request source physical address field within the address discovery request, in accordance with the determining step, by replacing the shared cluster address with a non-cluster mode physical address of the source host.

25 13. The method of claim 12 wherein the non-cluster mode physical address is a dedicated MAC address.

14. The method of claim 12 wherein the responding host comprises a network communication protocol-specific layer including an address discovery request handler that operates in the cluster mode and wherein the method further comprises passing, after the replacing step, the address discovery request to the address request handler, and
5 wherein the generating step comprises the further steps of:

first creating an initial address discovery response, by the address request handler, based upon the modified address discovery request, the address discovery response including:

10 the shared cluster address within the response source physical address field; and
the non-cluster mode physical address of the source host within a response target physical address field; and

15 second creating a revised address discovery response by:
first replacing the shared cluster address by the non-cluster mode physical address of the source host within the response source physical address field, and
second replacing the non-cluster mode physical address of the source host by the shared cluster address within the response target physical address field.

20 15. The method of claim 14 wherein the second creating step is performed by the responding host.

16. The method of claim 14 wherein the responding host comprises an address request handler that maintains an address resolution table including a set of entries
25 pairing network communication protocol-specific addresses with corresponding physical addresses, the method further comprising the step of:

storing, by the address request handler within the address resolution table, an entry including the non-cluster mode physical address and a corresponding network communication protocol-specific address of the source host.

17. The method of claim 14 wherein the network communication protocol-specific layer implements the TCP/IP protocol.

18. The method of claim 14 wherein the address discovery response further includes a response destination field specifying the shared cluster address assigned to the cluster further comprising rendering the address discovery response acceptable by the source host operating in the cluster mode by:

5 replacing, within the response destination field, the non-cluster mode physical address by the shared cluster address.

10

19. The method of claim 1 wherein the address discovery response further includes a response destination field specifying the shared cluster address assigned to the cluster.

15

20. A computer-readable medium containing computer-executable instructions for conducting physical address discovery, facilitating point-to-point communications between hosts of a cluster operating in a cluster mode wherein acceptable messages are addressed to a shared cluster address, the computer-executable
5 instructions facilitating performing the steps of:

receiving an address discovery request initiated by a source host seeking a physical address of a target host within the cluster; and

generating an address discovery response acceptable by the source host operating in the cluster mode and including:

10 a response source physical address field specifying a non-cluster mode physical address of the target host.

21. The computer-readable medium of claim 20 wherein the address discovery request is an ARP request.

15 22. The computer-readable medium of claim 20 wherein the non-cluster mode physical address of the target host is a dedicated address of the target host.

23. The computer-readable medium of claim 22 wherein the dedicated address
20 is derived from an IP address assigned to the target host.

24. The computer-readable medium of claim 20 wherein the non-cluster mode physical address of the target host is a shared address assigned to multiple hosts within the cluster.

25. The computer-readable medium of claim 20 wherein the generating step is executed in accordance with a further step of:
determining whether the address discovery request was issued by a source host operating in the cluster mode.

26. The computer-readable medium of claim 25 wherein the receiving, determining and generating steps are performed on a responding host that is distinct from an initiating host from which the address discovery request originates.

5 27. The computer-readable medium of claim 26 wherein the determining step comprises detecting that the address discovery request includes:

a request source physical address field specifying the shared cluster address assigned to the cluster; and
a request source network communication protocol-specific address field identifying a host within the cluster of hosts.

10 28. The computer-readable medium of claim 27 wherein the shared cluster address is a MAC address.

15 29. The computer-readable medium of claim 28 wherein the source network communication protocol-specific address field contains an IP address.

30. The computer-readable medium of claim 27 further comprising computer-executable instructions for performing the step of:

20 maintaining, by the responding host, a list of network communication protocol-specific addresses corresponding to hosts within the cluster.

31. The computer-readable medium of claim 27 further comprising computer-executable instructions for performing the steps of:

25 modifying the request source physical address field within the address discovery request, in accordance with the determining step, by replacing the shared cluster address with a non-cluster mode physical address of the source host.

30 32. The computer-readable medium of claim 31 wherein the non-cluster mode physical address is a dedicated MAC address.

33. The computer-readable medium of claim 31 wherein the responding host comprises a network communication protocol-specific layer including an address discovery request handler that operates in the cluster mode and wherein the method further comprises passing, after the replacing step, the address discovery request to the 5 address request handler, and wherein the generating step comprises the further steps of:

first creating an initial address discovery response, by the address request handler, based upon the modified address discovery request, the address discovery response including:

the shared cluster address within the response source physical address field; and
10 the non-cluster mode physical address of the source host within a response target physical address field; and

second creating a revised address discovery response by:
15 first replacing the shared cluster address by the non-cluster mode physical address of the source host within the response source physical address field, and
second replacing the non-cluster mode physical address of the source host by the shared cluster address within the response target physical address field.

20 34. The computer-readable medium of claim 33 wherein the second creating step is performed by the responding host.

25 35. The computer-readable medium of claim 33 wherein the responding host comprises an address request handler that maintains an address resolution table including a set of entries pairing network communication protocol-specific addresses with corresponding physical addresses, the method further comprising the step of:

storing, by the address request handler within the address resolution table, an entry including the non-cluster mode physical address and a corresponding network communication protocol-specific address of the source host.

36. The computer-readable medium of claim 33 wherein the network communication protocol-specific layer implements the TCP/IP protocol.

37. The computer-readable medium of claim 33 wherein the address discovery
5 response further includes a response destination field specifying the shared cluster address assigned to the cluster further comprising computer-executable instructions for rendering the address discovery response acceptable by the source host operating in the cluster mode by:

replacing, within the response destination field, the non-cluster mode physical
10 address by the shared cluster address.

38. The computer-readable medium of claim 20 wherein the address discovery response further includes a response destination field specifying the shared cluster address assigned to the cluster.

39. A host computer system including physical address discovery components facilitating point-to-point communications between hosts of a cluster operating in a cluster mode wherein acceptable messages are addressed to a shared cluster address, the computer system comprising:

- 5 a network interface for receiving an address discovery request initiated by a source host seeking a physical address of a target host within the cluster;
- a transport layer component for carrying out transport-protocol specific processing of network requests;
- intracluster address discovery logic interposed between the network interface and the transport layer component of the host system, the intracluster address discovery logic performing the step of:
 - 10 generating an address discovery response acceptable by the source host operating in the cluster mode and including:
 - 15 a response source physical address field specifying a non-cluster mode physical address of the target host.

40. The system of claim 39 wherein the address discovery request is an ARP request.

20 41. The system of claim 39 wherein the non-cluster mode physical address of the target host is a dedicated address of the target host.

42. The system of claim 41 wherein the dedicated address is derived from an IP address assigned to the target host.

25 43. The system of claim 39 wherein the generating step is executed in accordance with a further step of:

- determining whether the address discovery request was issued by a source host operating in the cluster mode.

44. The system of claim 43 wherein the determining step performed by the intracluster address discovery logic comprises detecting that the address discovery request includes:

- 5 a request source physical address field specifying the shared cluster address assigned to the cluster; and
 a request source network communication protocol-specific address field identifying a host within the cluster of hosts.

45. The system of claim 44 wherein the shared cluster address is a MAC
10 address.

46. The system of claim 45 wherein the source network communication protocol-specific address field contains an IP address.

15 47. The system of claim 44 further comprising a list of network communication protocol-specific addresses corresponding to hosts within the cluster.

48. The system of claim 44 wherein the intracluster address discovery logic comprises executable instructions for:

- 20 modifying the request source physical address field within the address discovery request, in accordance with the determining step, by replacing the shared cluster address with a non-cluster mode physical address of the source host.

49. The system of claim 48 wherein the non-cluster mode physical address is a
25 dedicated MAC address.

50. The system of claim 48 wherein the a transport layer component includes an address discovery request handler that operates in the cluster mode and wherein the method further comprises passing, after the replacing step, the address discovery request to the address request handler, and wherein the generating step performed by the intracluster address discovery logic comprises the further steps of:
- 5 first creating an initial address discovery response, by the address request handler, based upon the modified address discovery request, the address discovery response including:
- 10 the shared cluster address within the response source physical address field; and
- the non-cluster mode physical address of the source host within a response target physical address field; and
- second creating a revised address discovery response by:
- 15 first replacing the shared cluster address by the non-cluster mode physical address of the source host within the response source physical address field, and
- second replacing the non-cluster mode physical address of the source host by the shared cluster address within the response target physical address field.
- 20 51. The host of claim 50 further comprising an address request handler that maintains an address resolution table including a set of entries pairing network communication protocol-specific addresses with corresponding physical addresses, and wherein the address request handler stores, within the address resolution table, an entry including the non-cluster mode physical address and a corresponding network communication protocol-specific address of the source host.
- 25

52. The host of claim 50 wherein the a transport layer component implements the TCP/IP protocol.

53. The host of claim 44 wherein the address discovery response further includes a response destination field specifying the shared cluster address assigned to the cluster.

5 54. A method for processing point-to-point communications between hosts of a cluster operating in a cluster mode implemented by a network communication protocol-specific layer of each host, and wherein acceptable messages are addressed to a shared cluster address, the method comprising the steps of:

10 receiving an intracluster message issued by an initiating host including a non-cluster mode physical address of a target host within a message destination field; replacing, within the intracluster message by the target host, the non-cluster mode physical address by the shared cluster address; and presenting, after the replacing step, the intracluster message to the network communication protocol-specific layer.

15 55. The method of claim 54 wherein the replacing step is performed by a network load balancing component.

56. The method of claim 54 wherein the network communication protocol-specific layer implements TCP/IP.

20 57. The method of claim 54 further comprising the step of:
generating an intracluster response message including a non-cluster mode physical address for the initiating host within the message destination field.

58. The method of claim 57 further comprising the steps of:
receiving, by the initiating host, the intracluster response message including the
non-cluster mode physical address for the initiating host within the message destination
field;
5 replacing, within the intracluster response message by the initiating host, the non-
cluster mode physical address by the shared cluster address; and
presenting, after the replacing step, the intracluster message to the network
communication protocol-specific layer within the initiating host.
- 10 59. A method for performing point-to-point communications between hosts of
a cluster operating in a cluster mode implemented by a network communication protocol-
specific layer of each host, and wherein acceptable messages are addressed to a shared
cluster address, the method comprising the steps of:
receiving an address discovery request seeking a physical address of a target host
15 within the cluster;
determining that the address discovery request was issued by a source host
operating in the cluster mode;
generating an address discovery response acceptable by the source host operating
in the cluster mode and including:
20 a response source physical address field specifying a non-cluster mode
physical address of the target host;
receiving an intracluster message issued by the source host including a non-
cluster mode physical address of the target host within a message destination field;
replacing, within the intracluster message by the target host, the non-cluster mode
25 physical address by the shared cluster address; and
presenting, after the replacing step, the intracluster message to the network
communication protocol-specific layer.